13TH ANNUAL
MIT SLOAN
SPORTS ANALYTICS CONFERENCE
MARCH 1 - 2, 2019 | BOSTON, MA
presented by ESPN

Training Schedule Confounds the Relationship between Acute:Chronic Workload Ratio and Injury

A Causal Analysis in Professional Soccer and American Football[1]

Luke Bornn[2], Patrick Ward[3], Darcy Norman[4]

1. Introduction

In the past decade, significant efforts have been made to understand injury risk in sport using subjective (i.e. rating of perceived exertion) and objective (i.e. inertial sensor outputs) player-monitoring strategies (Halson, 2014; Bourdon et al., 2017). One metric in particular that has received significant attention is the acute:chronic workload ratio (ACWR), defined at time t as

$$ACWR_t = L^7{}_t \,/\, L^{28}{}_t$$

where $L^7{}_t$ and $L^{28}{}_t$ are the cumulative player workloads in the last 7 and 28 days, respectively (Hulin et al., 2014; Hulin et al, 2015). In the past 5 years, numerous academic papers across multiple sports have concluded that ACWR is predictive of injury risk, and as a result the ACWR has become standard practice in professional sports to manage player workloads (Malone et al., 2017a; Malone et al., 2017b; Colby et al., 2017).

In this paper, we use schedule and training data from two professional teams in Serie A (Italian soccer) and the NFL (American football) to show that training schedule confounds the ACWR-injury relationship, leading to overestimates of the magnitude of the relationship and calling into question earlier studies which have ignored this confounder.

2. Data

Our data consists of daily training loads over one season for a team in each of Italian soccer and American football. In both cases the data is collected from inertial sensors and combined with anthropometric measurements to get a notion of total physical load experienced by each player per session. While we focus on a single metric (player load) to simplify presentation, note that the confounding we observe in the paper applies across the suite of metrics most teams collect, such as acceleration and sprint counts, as well as subjective measures such as rating of perceived exertion. We remove goalkeepers and quarterbacks due to the unique physical demands of those positions, and average the player loads per session to obtain a daily average team-wide training load.
We calculate acute load as a 7-day average leading into the session, and chronic load as the corresponding 28-day average. Illustrating in words, ACWR on October 29 is calculated as “sum of player load from Oct 22-28 divided by 7” divided by “sum of player load from October 1-28 divided

[1] To encourage practitioners to dive into this work, code and representative pseudo-data is included at the first author’s Github profile.
[2] Sacramento Kings
[3] Seattle Seahawks
[4] US Men’s National Soccer Team

by 28." Load is assumed to be zero on non-training days. In Figure 1 we show an image of the raw data across the full season, along with acute (red) and chronic (blue) loads.

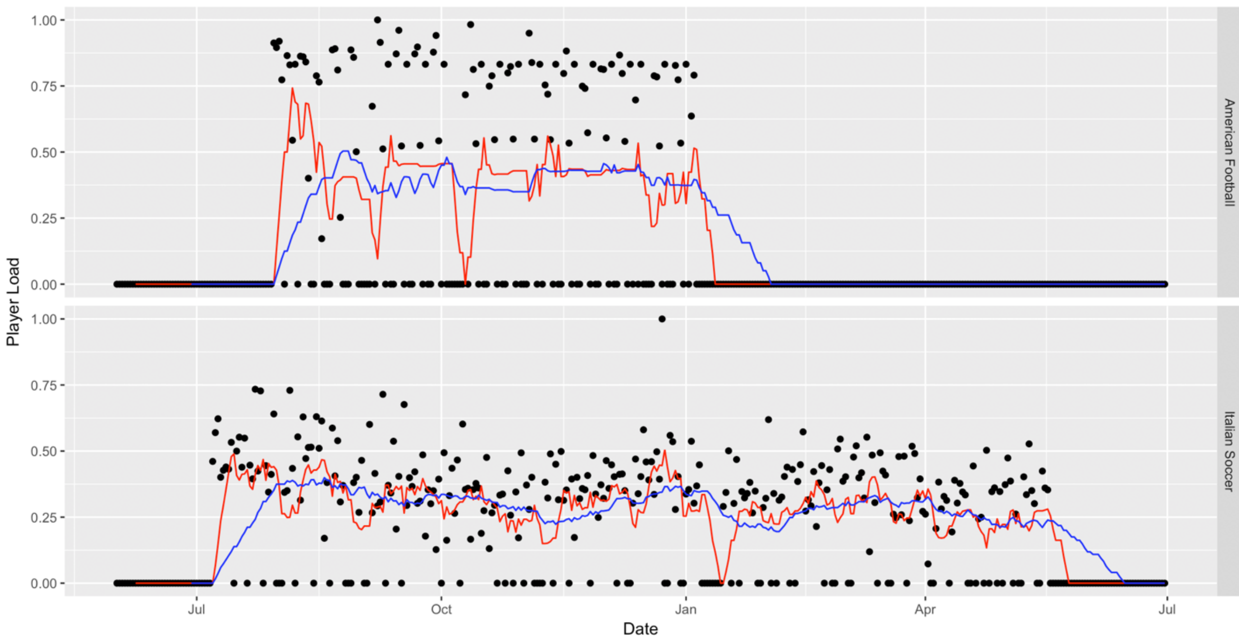

Figure 1. Raw data showing normalized loads in each sport, as well as acute (red) and chronic (blue) loads through the season.

3. A Null Model of Injury

The fundamental observation of this paper is that even if we simulate season-long training schedules (and injuries) such that injuries only depend on the player load in the current session (and hence ACWR has no influence), when we study these simulations at season's end ACWR will be significantly correlated with injury. Graphically interpreted, the standard approach to understanding the ACWR-injury relationship (Figure 2, left) ignores confounding schedule effects Φ (Figure 2, right), and hence mis-captures the relationship.

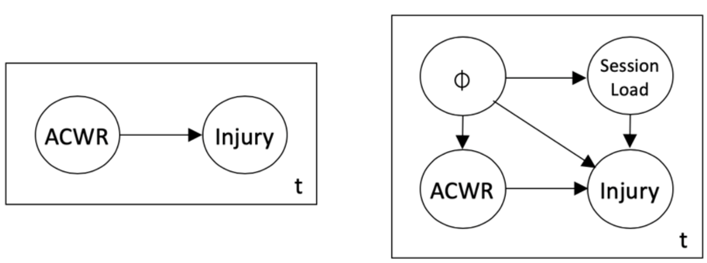

Figure 2. Diagram of ACWR-injury relationship ignoring the confounding of schedule (left) and accounting for scheduling effects Φ (right).

To study this confounding, we conduct a large-scale simulation driven by the underlying training loads collected from the aforementioned teams. Specifically, we simulate 1000 player-seasons using the observed team schedules and training loads, whereby the probability of injury in a given session is directly proportional to player load in that session. We calibrate this to induce injuries in approximately 5% of training sessions, though the paper's results are robust to higher or lower injury rates. By simulating in this way, we are specifying that elements of training prior to the session are irrelevant to a player's likelihood of being injured that session.

Next, we define the acute:chronic flag ("AC flag") as any session where the ACWR falls outside of the (0.8, 1.3) range (Gabbett, 2016). After simulating the 1000 seasons, we find a statistically significant difference in AC flag prevalence between injury and non-injury sessions (soccer: $p = 2.6 \times 10^{-12}$; football: $p = 1.4 \times 10^{-7}$). Thus even though we simulated seasons in such a way that ACWR had no impact on injury risk, a post-hoc study finds a meaningful relationship between ACWR and injury! In fact, the presence of AC flag prior to a session is predicted to result in a 6.8% and 10.5% increase in injury risk in Serie A and the NFL, respectively, even though by construction both should be 0%.

There has been a tremendous amount of work to date studying the ACWR-Injury relationship in a manner akin to the above, the large majority finding similar levels of statistical significance. This has led to the broad adoption of ACWR as a diagnostic tool for injury across all levels of sport. While our work does not falsify these studies, it does raise the possibility that the significance found in these papers could be the result of confounding of schedule.

4. Correcting for Schedule and Injury Session Load

The underlying thesis of this paper is that due to the unique nature of training loads across sports (training camps, international breaks, etc.), the ACWR-injury relationship is confounded with schedule. More specifically, ACWR is correlated with injury through the current session load. Figure 4 demonstrates this effect in both American football and Italian soccer, where we see that sessions with an AC flag leading into the session have higher loads than sessions without the flag.

We now explore whether conditioning on the current session load mitigates the confounding of schedule. Specifically, we fit two logistic regression models with injury as an outcome, the first using just AC flag as a covariate, and the second also including session load. The initial model results in the same p-values as in the baseline tests listed above, as expected. However, after conditioning on session load this relationship disappears (soccer: $p = 0.33$; football: $p = 0.70$).

While conditioning on the current session load removes the confounding effects, this is a unique possibility created (by design) in the simulation environment. In practice, the data from the injured player is often censored or missing completely, as treatment of the player takes precedence over data collection. This censoring suggests an alternative method for adjusting for this confounding, namely survival models (Cox, 2018). When the injured player's data is available up until the moment of injury, survival models may be applied directly. Specifically, one would include the player load until the moment of injury as the time to injury, and every non-injured player would be censored at their full observed player load. In other words, healthy players are considered censored before an injury was observed. In fact, applying this technique to our data by randomly simulating injury times within sessions, the relationship between AC flag and injury indeed becomes non-significant (soccer: p = 0.22; football: p = 0.41).

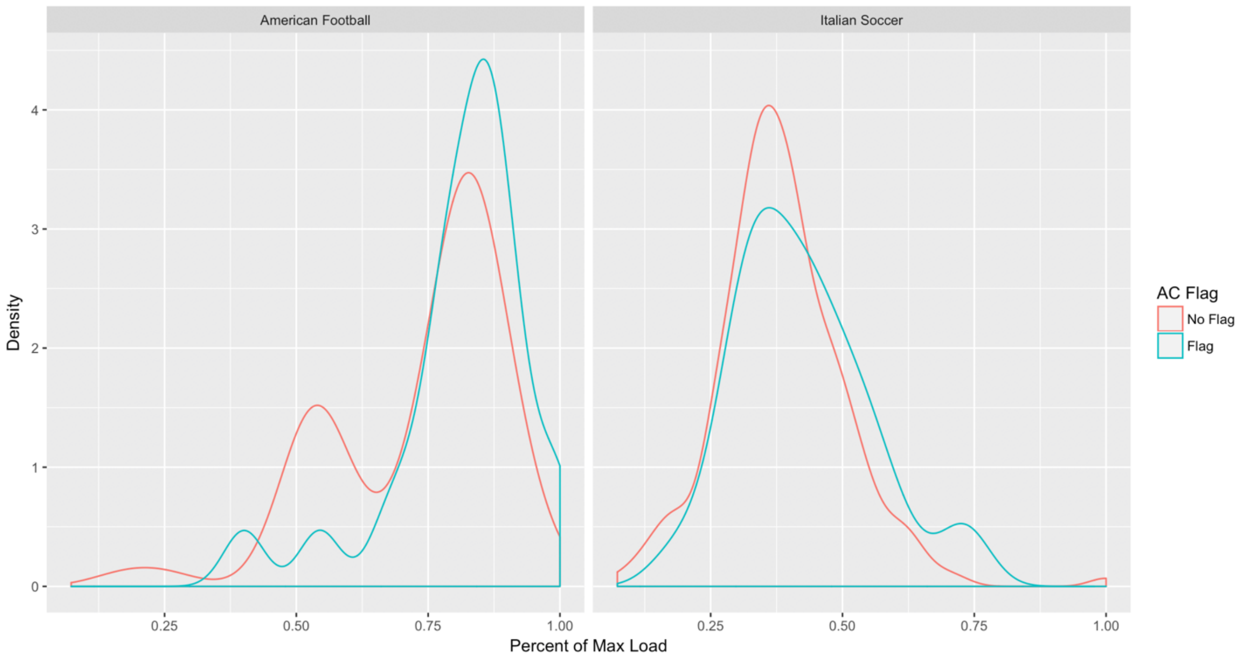

Figure 4. Distributions of load differ between flagged sessions and non-flagged sessions, indicating that schedule is confounding the ACWR-Load relationship, and hence also the ACWR-Injury relationship.

When an injured player's data is not available, we may employ the information from his teammates to understand injury at the team or position level. As an example, Ward et al. (2018) use position averages to show the relationship between session load and injury in American football at the position level. In that case, the underlying estimand changes from an individual's injury risk, to the likelihood of an injury within a given position group. Alternatively, if the focus is at the individual level, one could employ survival models alongside latent variable models, imputing or inferring the injured player's missing load data. Such approaches require significant knowledge of Bayesian hierarchical models and causal inference, however. Interested readers are referred to Gelman et al. (2013) and Imbens and Rubin (2015) for a broad treatment of such approaches.

5. Conclusions

In summary, we demonstrate that causal conclusions about the ACWR-injury relationship are prone to confounding from schedule. We use Monte Carlo methods combined with training load data from two sports to illustrate the effect that the yearly training calendar has on the ACWR-injury relationship. We then propose options to mediate this confounding. Our study impacts not only the academic discourse around the ACWR, but also gives practitioners a more realistic expectation of its value in predicting injury.

6. References

[1] Halson SL (2014). Monitoring training load to understand fatigue in athletes. Sports Med; 44, Suppl 2: 139-147.

[2] Bourdon, PC. Cardinale, M. Murray, A. Gastin, P. et al. (2017). Monitoring athlete training loads: Consensus statement. Int J Sports Physiol Perf; 12(S2): 161-170.

[3] Hulin, BT. Gabbett, TJ. Blanch, P. Chapman, D. Orchard, JW. (2014). Spikes in acute workload are associated with increased injury risk in elite cricket fast bowlers. Br J Sports Med; 48(8): 708-712.

[4] Hulin, BT. Gabbett, TJ. Lawson, DW. Caputi, P. Sampson, JA. (2015). The acute:chronic workload ratio predicts injury: High chronic workload may decrease injury risk in elite rugby league players. Br J Sports Med; 50(4): 213-236.

[5] Malone, S. Owen, A. Mendes, B. Collins, KD. Gabbett, TJ. (2017). The acute:chronic workload ratio in relation to injury risk in professional soccer. J Sci Med Sport; 20(6): 561-565.

[6] Malone, S. Roe, M. Doran, DA. Gabbett, TJ. Collins, K. (2017). High chronic training loads and exposure to bouts of maximal velocity running reduce injury risk in elite Gaelic football. J Sci Med Sport; 20(3): 250-254.

[7] Colby, MJ. Dawson, B. Peeling, P. Heasman, J. Rogalski, B. Drew, MK. Stares, J. Zouhal, H. Lester, L. (2017). Multivariate modelling of subjective and objective monitoring data improve the detection of non-contact injury risk in elite Australian footballers. J Sci Med Sport; 20(12): 1068-1074

[8] Cox, David Roxbee. *Analysis of survival data*. Routledge, 2018.

[9] Lolli, L. Batterham, AM. Hawkins, R. Kelly, DM. Strudwick, AJ. Thorpe, R. Gregson, W. Atkinson, G. (2017). Mathematical coupling causes spurious correlation within the conventional acute-to-chronic workload ratio calculations. Br J Sports Med. E published ahead of print.

[10] Gabbett, TJ. (2016). The training-injury prevention paradox: Should athletes be training smarter *and* harder? Br J Sports Med; 50(5): 273-280.

[11] Ward, P., Tankovich, M., Ramsden, S., Drust, B., Bornn, L. (2018). Volume and Intensity are Important Training Related Factors in Injury Incidence in American Football Athletes. *Sloan Sports Analytics Conference 2018.*

[12] Gelman, A., Stern, H. S., Carlin, J. B., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis*. Chapman and Hall/CRC.

[13] Imbens, Guido W., and Donald B. Rubin. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press, 2015