An enhanced metric of injury risk utilizing Artificial Intelligence

Calham Dower, Abdul Rafehi, Jason Weber, Razali Mohamad / Alerte Digital Sport / calhamd@alerteds.com

The value of an effective injury risk management system can be easily quantified. The average Australian Rules footballer earns $15,000 per game. For each game a player is unavailable, the team loses $15,000 in performance value. If a system can prevent even 10 games lost through injury a year (the equivalent of 3 injuries), this can mitigate value lost by $150,000. Clearly, such a system will provide considerable return on investment for a team. This is exacerbated when considering top end talent ($25,000+ per game), and in particular other sporting leagues, such as the EPL ($100,000+ per game).

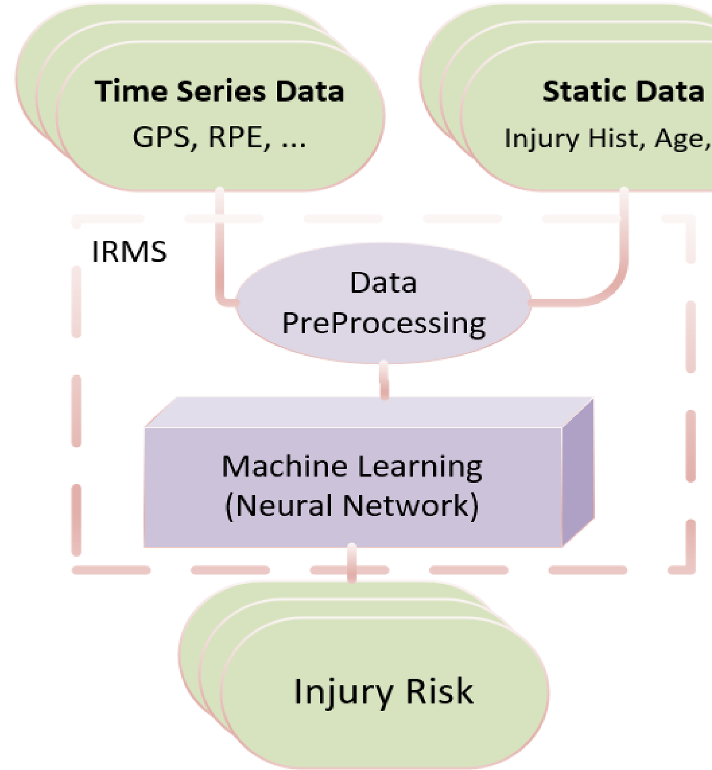

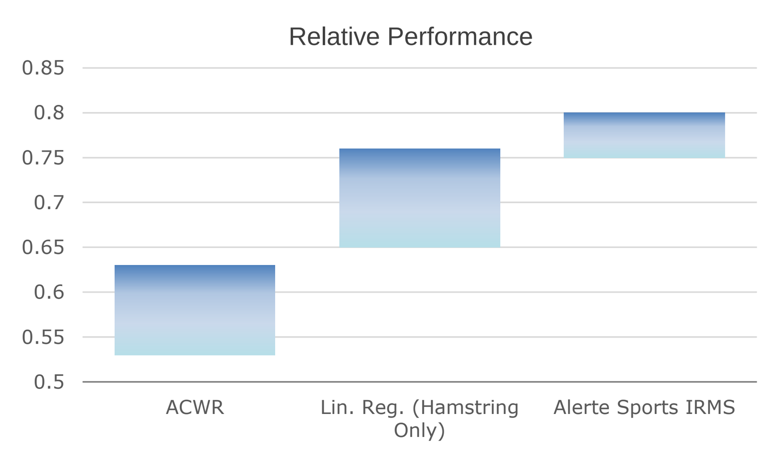

Methodology and Results

The methodology utilizes a combination of factors, including GPS volume metrics, injury history information and expected session intensities to calculate an injury risk factor. It uses a combination of features that convey more information than any one feature in isolation. These features are pre-processed to maximize information without leading to an excess of inputs into the system, ensuring that it generalizes to work on unforeseen data.

The neural-network learns the features that are present in injury scenarios and learns to recognize them in future events.

AI? But Why?

IRMS uses ANNs as its modelling paradigm. With enough degrees of freedom, hidden layers and nodes, an ANN is able to approximate any non-linear piecewise function. This allows the system to be able to model very complex systems and form disjointed decision boundaries in a large dimension space. The Artificial Neural Network, if trained correctly, is able to discover patterns in the underlying signal that lead to increased injury risk. Analogous to this would be a human recognizing a cat because it has ears and whiskers, rather than having seen EXACTLY that cat before.

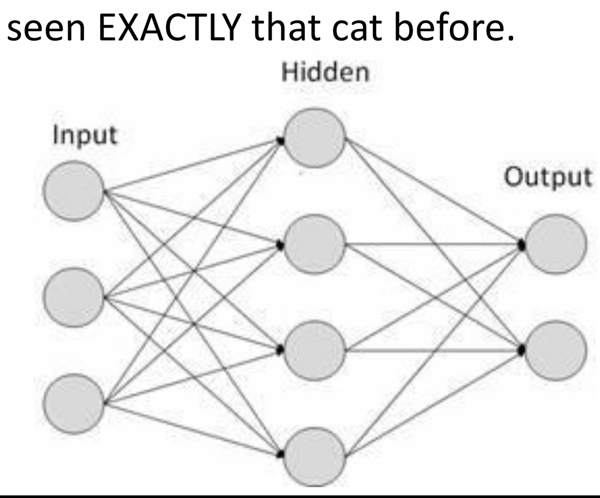

Common Problems and Performance

There are many common pitfalls associated with attempting injury risk estimation and its useful application in practice. These lead to misleading performance rates that are not realistically achievable.

Final day hidden bias

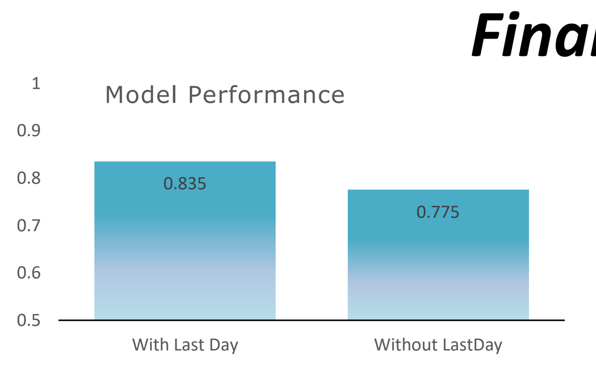

Assume a player normally runs 10km on gameday. If the player is injured mid game, this may read as 4km. Using this as input is erroneous, as the shortened session is an outcome of the injury, rather than an influence on it. This erroneous behavior is often included in similar studies, and can easily skew results and real-world performance.

Does the model mimic real-world expectations

Players are much more likely to be injured on gameday, and therefore it is completely reasonable to expect any injury prediction model to have elevated likelihood of injury on gameday. The model, without being told which day is gameday, automatically assigns higher risk on gameday.

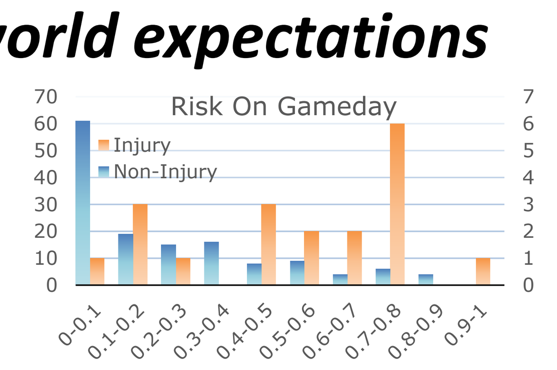

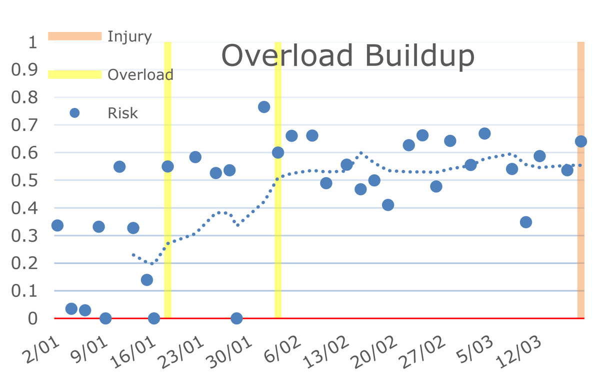

We can see that injury risk periods are aligned with periods of overload. When overload events (elevated DOMs and other soreness) are included in training, we see that these often lead into injury events, and as such increase the predictive accuracy of the models. The models have learned this and often increase injury risk in the presence of repeated overload events as one might intuitively expect.

Over the course of one season, individuals generally enter and fall out of periods of elevated injury risk. By applying a moving average filter across the risk scores, it is seen that a player enters periods of exacerbated risk. Observing these periods is critical in player management, and can assist teams in managing players to maximise performance while minimising injury risk

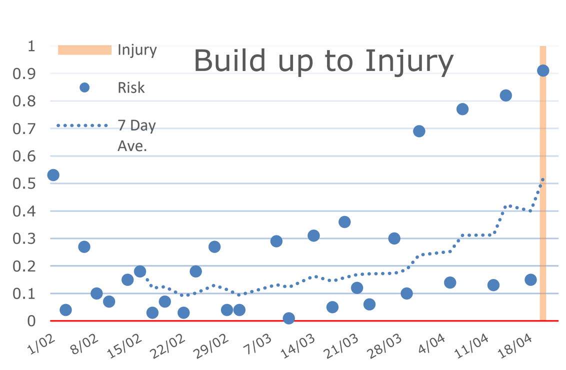

Does performance improve over time?

A model's ability to correctly predict an outcome is linked with the amount of data that model as been trained on. As more data is utilized, we expect the model to better generalise about the factors that lead to elevated risk, and thus become a more accurate predictor.

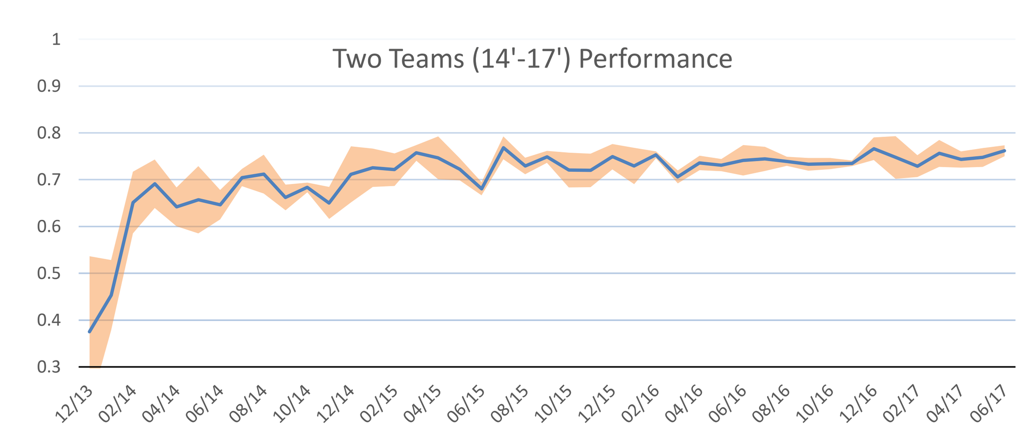

The usability of a model is tied to the amount of data required for it to converge. If it takes years of individual team data before a model provides useable results, then clearly such a model is of little value in a real-world environment.

Does utilizing multiple datasets improve performance?

A key benefit of our system is the ability to leverage, in an anonymous way, data from different teams. While we know it might take 2 or 3 seasons to tailor a useable model for a team from scratch, we show that this process can be expedited by using data from a different team to kick-start a model.

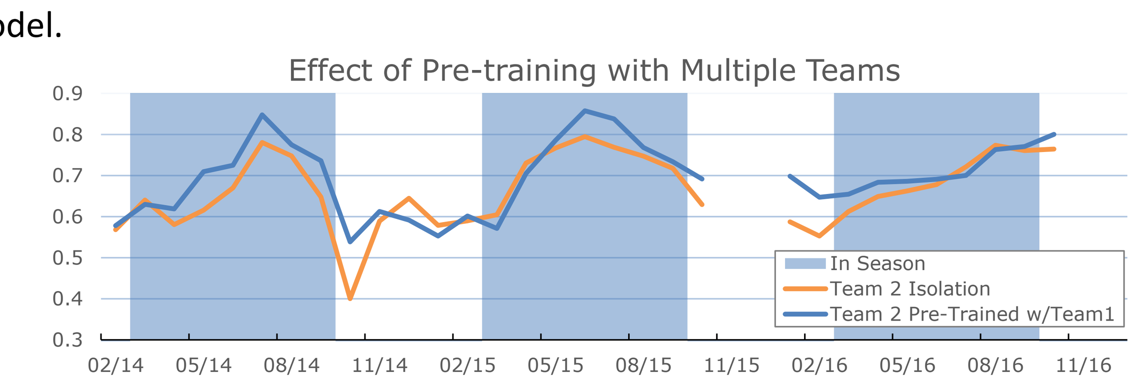

Conclusion

- As the output is a continuous score between 0 and 1, the result can be utilized by practitioners as a risk rating
- The system is shown to improve over time which indicates it has not yet reached its peak potential
- The system is robust to common problems associated with predicting injury risk.
- The system is accurate enough to provide insights to Athlete performance professionals.