Big data may not be enough to predict injuries in the NFL

The NFL is hoping big data tools can help bring down the number of concussions, ligament tears, and other injuries sustained in each game of professional football. Currently, the injury count per game is holding steady at an average of six or seven. League engineers are working with Amazon Web Services to apply machine learning and artificial intelligence tools to player data, with the hope of finding in-game situations that commonly lead to injury, *The Wall Street Journal* reported this week.

"Ultimately, we will be able to identify injury risk scenarios, and we will be able to predict injury risk scenarios, and we will be able to find innovations that will make the game safer for our athletes while maintaining high quality of play," Jeff Crandall, chair of the NFL's engineering committee, said during the program announcement.

The NFL and Amazon have vast resources at their disposal. But injuries, especially in chaotic sports like football, are incredibly hard to predict. "It's the holy grail. Everyone wants to do it, and no one can," says Zachary Binney, an epidemiologist and consultant who has worked with Major League Baseball and college sports teams on injury prevention. "I'm skeptical until I see results."

Predicting injuries is challenging because there are so many factors that could contribute to a possible injury, from an athletes physical characteristics on a particular day to slight divots on a field. One athlete might have five attributes that research shows puts them at risk for an injury and still not get hurt, but another might look perfectly fine

and tear a ligament the next day. “It is just an incredibly difficult problem,” Binney says.

The Amazon Web Services partnership will try to close the gap with league-level data from the NFL’s “Next Gen” stats, which capture location data, speed, and acceleration for every player on the field hundreds of times a minute through microchips in their pads. It also includes video footage of games, information on playing surface and environmental factors, and anonymized player injury data, according to the NFL. It doesn’t collect data on how hard body parts are hitting the ground or other players, which is one limitation, Binney says. But it can see, with granular detail, how and at what speed a player ran a play, changed direction, or made a tackle. The goal is to find out if any common elements of football are more likely than others to lead to any injury.

“You could look at what happens when a wide receiver moving this quickly makes this sharp turn, and might be able to tease something out,” Binney says.

The league-level data only includes some measures of player activity. Individual teams have more granular level data on their players, typically tracking things like heart rate, fatigue, hydration, and other measures — all of which can contribute to injury risk for a particular player. Other risk factors for injuries in football include flexibility, injury history, strength, and body composition. However, much of the player-specific data stays at the team level, to avoid giving their opponents potentially useful information about how their players are doing.

Player health data will not be included in the injury prediction program, according to an email from an NFL spokesperson to *The Verge.* That might affect its predictive power. “It will be really interesting to see. I don’t know what impact that might have, and I can’t imagine that they do yet either,” Binney says.

Even without the more granular information, the league has data on the athletes from all 32 teams, which gives them more to work with. “You’re losing some of the resolution on the data, but increasing the sample size,” Binney says.

In the past, employees at the NFL have manually gone over hundreds of hours of game footage and helmet impacts to identify situations that lead to injury, and made changes — like updates to the kickoff rules — with the goal of preventing them. Binney speculates that the project could lead to additional changes, but that any information they’re able to gather could have additional value beyond that. “One thing they could do is put the information out there and tell coaches that when they ask a lineman to do some kind of block, or route, it creates the sorts of changes in direction or deceleration when we see bad things happen,” he says. It’s in the best interest of coaches, after all, to keep players healthy, and there could be alternatives to riskier plays.

If the NFL’s efforts into injury prediction and prevention prove effective, they might offer a road map to other sports, as well. Binney says it’s a positive step. “I’m excited to see it happen, even though I’m cautious about how much we might be able to draw from it.”

https://www.theverge.com/2019/12/6/20999403/amazon-nfl-injuries-concussions-big-data-machine-learning