New data shows baseball managers when to replace the starting pitcher

February 28 2014, by Peter Dizikes

Credit: Thinkstock

Last October, the Detroit Tigers won the first game of the American League Championship Series against the Boston Red Sox; the Tigers led the second game, 5-1, going into the eighth inning in Boston's Fenway Park, with one of the league's best starting pitchers, Max Scherzer, on the mound. They were six outs from taking command of the series.

Then Tigers manager Jim Leyland made a disastrous decision: He turned the game over to his bullpen, which promptly blew the lead in the eighth inning and lost the game in the ninth inning. Instead of the Tigers holding a 2-0 series lead heading back to their own ballpark for three games, the series was tied, 1-1, and the Red Sox went on to win it in six games.

Should Leyland have taken Scherzer out of the game? No, according to a unique model built by two MIT computer scientists, which indicates that major-league baseball managers have significant room to improve their decision-making.

Indeed, while managers sometimes seem to remove starting pitchers too hastily, as in Scherzer's case, they even more frequently stick with starting pitchers too long: The study finds that from the fifth inning on, in close games, pitchers who were left in games when the model recommended replacing them allowed runs 60 percent of the time, compared to 43 percent of the time overall.

"Clearly the most important decision a manager makes is changing pitchers," says John Guttag, the Dugald C. Jackson Professor of Computer Science and Engineering at MIT. In making those decisions, he adds, "I think there's definitely room for improvement."

Guttag developed the model with Ganeshapillai Gartheeban, one of his PhD students in MIT's Computer Science and Artificial Intelligence Laboratory. Their paper, "A Data-driven Method for In-game Decision Making in MLB," is one of eight finalists in the research paper competition at this year's MIT Sloan Sports Analytics Conference (SSAC), being held today and tomorrow at the Hynes Convention Center in Boston.

To conduct the study, Guttag and Gartheeban took data from the 2006

through 2010 major-league baseball seasons. They used the first 80 percent of the games in those seasons to build a model of how pitchers fare over the course of a game, concentrating on Pitcher's Total Bases (PTB)—an aggregate measure of hits and unintentional walks allowed—as the leading indicator of future performance. PTB, they note, is a more granular measure of pitcher performance than runs allowed.

"There is a lot of randomness involved in giving up a run," Gartheeban observes. "If you train a model on that, it will attach itself to noisy patterns. We go below that, to a more fundamental variable."

The researchers then tested the model on the final 20 percent of those seasons. Over 21,538 innings, the Guttag-Gartheeban model disagreed with the manager's decision regarding his starting pitcher 48 percent of the time. About 43 percent of the time, the manager left the starting pitcher in when the model indicated he should be replaced. In just 5 percent of the cases did managers pull starting pitchers when the model suggested they should stay in the game—the scenario from the Tigers-Red Sox game last fall.

Admittedly, in those latter cases, "there is no way to know how the starter would have done had he not been removed," as the paper notes.

By focusing on in-game decision-making, the paper brings to baseball a subject that has proven popular in football—where many studies have shown that teams should go for it on fourth down, rather than kicking. Despite the data, NFL coaches have been slow to change their ways.

Guttag—a lifelong New York Yankees fan—hopes big-league managers will be quicker to use this kind of data, although he emphasizes that they have to consider many complicating factors: how rested the bullpen is, the upcoming schedule of games, and more.

"The managers are considering a lot of things," Guttag acknowledges. "I wouldn't come to the conclusion that the entire gap [between the model and the actual decisions] is due to managers making bad decisions. The managers may well be making better decisions [in some cases] than we would if we knew all the things they have to consider."

Channeling the data deluge

Founded in 2007 and originally held on MIT's campus, SSAC has since grown to become the biggest, most prominent event of its kind in global sports. This year's session includes more than 30 panel discussions, featuring the likes of NBA Commissioner Adam Silver, retired basketball coach Phil Jackson, and Red Sox owner John Henry, among dozens of other prominent coaches, general managers, players, and analysts.

The conference's research paper competition—whose winner earns $20,000—features multiple entries based on the new optical tracking data now being gathered in all 30 NBA arenas. The SportVU system, as it is known, records the coordinates of all players, officials, and the ball, at 25 frames per second.

Using the SportVU data, Guttag, along with students Jenna Wiens and Armand McQueen, has produced another finalist entry in the SSAC research paper competition, titled "Automatically Recognizing On-Ball Screens." Like the baseball paper, this one applies machine-learning techniques to a wealth of information.

In this case, Guttag, Wiens, and McQueen have developed a machine learning-based system for recognizing when an on-ball screen—also known as the pick-and-roll, an essential part of basketball's offensive flow—occurs in the SportVU data. That could allow coaches and scouts to comb through hours of footage more efficiently.

"In the modern NBA, on-ball screens are used as really the heart of offenses," says McQueen, an MIT junior with a strong interest in sports analytics. In a trial run of the technique on 14 NBA games, the researchers found, their system identified 80 percent of on-ball screens with a confidence level of 82 percent.

Of the teams involved, the researchers also found that the specific screening patterns used by the Golden State Warriors and the Houston Rockets were most similar to each other, the kind of affinity that could provide more assistance to coaches and scouts.

"It's the first step in a machine-learning, pattern-recognition system," says Wiens, who will receive her PhD in electrical engineering and computer science this spring and is applying for academic jobs as a computer science professor. "I think this work is really a proof of concept of what can be done in the NBA, and it's really the tip of the iceberg."

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: New data shows baseball managers when to replace the starting pitcher (2014, February 28) retrieved 25 April 2024 from https://phys.org/news/2014-02-baseball-pitcher.html

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.