Industrial Systems Engineering and Management

Trust, but verify: a new framework for smarter decision-making

Historical data can guide or mislead. A new theoretical framework shows how to lean on it, and when to walk away.

Say you're planning to launch a product in Singapore and have a sea of historical sales data — but from the United States. Would you leverage this data to guide your decisions? Or would that be too speculative, given the many ways in which the two markets differ?

The conundrum of knowing when to trust past data, and when to ignore it, has long been a pain point for decision-makers in business, healthcare, technology and more. A theoretical study by Assistant Professor Cheung Wang Chi from

Graphic images in this newsletter were generated using AI and intended only as a visualisation of general concepts or ideas related to the research.

the Department of Industrial Systems Engineering and Management, College of Design and Engineering, National University of Singapore, offers some pointers. His work, published in the *Proceedings of the 41st International Conference on Machine Learning*, looks into the thorny problem of learning and optimising decisions in the moment, while cautiously sifting through the influence of past experience.

"At its core, this is a mathematical take on a very human dilemma: how to act judiciously in the present, when your only guide might be flawed memories of the past," says Asst Prof Cheung.

Assistant Professor Cheung Wang Chi and his PhD student devised a new theoretical framework to lay the foundation for future algorithms that make more reliable decisions in uncertain environments.

To trust, or not to trust?

To frame the problem, Asst Prof Cheung turned to a model inspired by probability theory, sometimes referred to as the "multi-armed bandit" problem. Picture a row of slot machines ("one-armed bandits") in a casino. Each machine gives different, unknown payouts. You want to find the one that pays the most, but you don't know which it is. So, you try each one, keep track of the results, and gradually favour the better performers.

Now imagine you're allowed to peek at records from past players — someone else's experience with the machines. There's a catch: those records may be inaccurate. Perhaps the machines were refurbished. Maybe you're not even in the same casino. Challenges abound. How do you decide whether that past data still applies — and whether using it will help or hinder you?

Asst Prof Cheung's work explores a pure theory setting — there's no real-world dataset, no simulations of consumer behaviour or ad clicks. Instead, it's about the mathematics behind decision-making — what you can and cannot do with imperfect prior knowledge.

One of the study most important results is a negative one. It demonstrates that even if you have access to historical data, it won't always help, and in some cases, it can't help at all. Unless you have some additional clue about how different the past and present might be, your safest bet is to ignore the old data entirely and proceed as if starting from scratch. "We initially thought it might be possible to create an algorithm that always knows when to trust the past," says Asst Prof

Cheung. "But our failed attempts convinced us otherwise, and that led to our eventual impossibility result."

To get around this limitation, Asst Prof Cheung and PhD student Lyu Lixing proposed a new algorithm, MIN-UCB, which makes use of a rough estimate of the "margin of error" between past and present. If this margin is small (for example, a 5% shift in customer behaviour between two seasons) the algorithm cautiously incorporates past data. If not, it plays it safe. "The algorithm compares two estimates of an option's potential — one based on both past and present, the other on current data alone — and choose the conservative of the two," says Mr Lyu.

In practice, this could apply to sales decisions where customer demographics or preferences have shifted slightly. For example, if an e-commerce platform notices a modest increase in new users, or if it's applying past purchasing data from one country to another with broadly similar market behaviour. In these cases, the algorithm can still make good use of history, without being blindly swayed by it.

This mechanism makes the algorithm highly adaptable. When historical data closely resembles current patterns, it improves on traditional approaches. When the data proves unreliable, it falls back to safer strategies. The team shows that this adaptability is not only intuitive but optimal within their theoretical framework.

A more discerning AI

"Our study raises some broader, and rather philosophical, questions," says Asst Prof Cheung. "How should we draw lessons from the past? What does it mean to 'know' something in a world that keeps changing?"

"Our study raises some broader, and rather philosophical, questions."

The researchers' work provides a foundation for future algorithms that can make more reliable decisions in uncertain environments — where history may inform, but not dictate, the present. As for next steps, the work is already expanding into more complicated domains. For instance, Asst Prof Cheung is exploring its application to reinforcement learning and online resource allocation, including a case study involving coupon assignments at an e-commerce company in Indonesia.

"In a world awash with data, knowing when not to trust information is just as important as knowing when to rely on it," adds Asst Prof Cheung. ◆