How Infusion Scheduling is Like a Game of Tetris

BY MOHAN GIRIDHARADAS

If you're of a certain age, you remember Tetris, a tile-matching puzzle video game originally designed and programmed by Russian game designer Alexey Pajitnov. Back in 1984—before "going viral" was a term—millions of people became obsessed (some might even say addicted) to this simple game, which challenged players to stack blocks of different shapes and sizes as efficiently as possible. Stacked improperly, the blocks would reach to the top of the screen—game over. Stacked optimally, the game could go on for hours.

Tetris & Infusion Scheduling

It turns out scheduling patients with varying treatment lengths is a lot like Tetris, only the odds of winning are stacked much more against you for a variety of reasons, including:

- Most infusion centers schedule patients on a first-come, first-served basis without an understanding of the overall "portfolio of patients who will be in adjacent chairs" at the time of the proposed appointment slot being planned.
- Between the randomized arrangement of appointments, late arrivals, delays in the lab and/or pharmacy, patients who experience an adverse reaction, nurses calling in sick, and the clinics running late, it is almost inevitable that there is a mid-day "crunch" sometime between 10:00 am and 2:00 pm when all of the chairs are full with numerous people in the waiting room.
- Simple spreadsheets or traditional EHR (electronic health record) approaches are not designed to create an optimal solution for scheduling infusion appointments.

By the Numbers

The complexity of the math required to create an optimized daily schedule is daunting. As an example:

- Take a 35-chair infusion center that operates from 7:00 am until 7:00 pm each day treating five types of appointments: 1 hour, 2 hours, 3-5 hours, 6-8 hours, or 9 or more hours.
- Assume that there is a sufficient number of nurses available in order to simultaneously seat 4 patients at a time to start their treatment at 10-minute intervals. That's 256 possible start times or "slots" per day (64 start times 7:00 am, 7:10 am, 7:20 am etc., until 5:30 pm multiplied by 4 patients at each start time, or 64 X 4 = 256)
- The number of possible ways these patient appointments can be arranged is a number with more than 100 zeros behind it. To put such a gigantic number into perspective, if we were to put all of the water in all of the oceans into gallon jugs, the number of jugs needed would be a number with only 40 zeroes behind it.

So, if the numbers are stacked against cancer center schedulers, how can they possibly arrange the roster of appointments—not just for today, but for every day for the next several weeks—in a way that allows the cancer program to keep up with increasing patient volumes, prevent excessive wait times, and keep operational costs down?

The schedule must also deal with multiple operational constraints, including the number of nurses, the number of chairs, the hours of operation, and constraints imposed by related services, such as the lab or the pharmacy—the list goes on.

As if that scenario isn't sufficiently challenging, infusion centers have to deal with the consequences of expected and unexpected variability—very few appointments in an infusion center will start and end exactly at the planned time. Therefore, any schedule has to be resilient to the inevitable shocks that will occur.

Making the Numbers Work for You

To end up with an answer that works efficiently across multiple constrained resources, you need to start with a sophisticated prediction model to accurately estimate the total number of patients for each day of the week and to have a realistic assessment of the mix of treatment durations (i.e., how many 1-hour, 2-hour, 3-hour appointments are likely on a Monday or a Tuesday).

Armed with an accurate prediction, the next step is to optimize the scheduling template based on the multiple operational constraints that are relevant for the specific center. Unfortunately, in a problem with multiple resource constraints, most solutions are suboptimal and are likely to introduce major bottlenecks at different times of the day. Building a schedule that balances the workload (acuity, number of patients supported, etc.) across nurses may work well for the nurses but may create situations where the center runs out of chairs in which to seat patients. Similarly, attempting to schedule to available chairs makes it likely that the bottleneck will be created as a consequence of not having enough nurses to treat patients at certain times of the day.

And what about pods? Splitting the problem into pods makes it easier to conceptualize the demand on multiple resources, but creating smaller groupings of nurse and chair resources limits the overall efficiency, and locks you into patterns that may not work once the variability of the day hits—with some patients arriving late, others who have a bad reaction to a drug and need to stay in a chair longer than planned, and still others who need to be urgently added to the schedule.

The key ingredient, of course, is data, more specifically EHR data. Inspired by the likes of Toyota and just-in-time Lean manufacturing practices, data science and mathematics are changing the face of healthcare scheduling, making it possible to optimize healthcare operations in ways that have not been done before. For example, LeanTaaS data scientists mine scheduling patterns and create optimal templates that are customized to an infusion center and automatically eliminate the mid-day peak by flattening the chair utilization profile throughout the day on every single day. These templates then incorporate machine learning algorithms to continuously improve thereby adapting to the changing volumes, mix and provider patterns.

This mathematical approach to infusion center scheduling is already delivering impressive results. Providers like Stanford Health Care, UCHealth, NewYork-Presbyterian, the UCSF Helen Diller Family Comprehensive Cancer Center, the Huntsman Cancer Institute, Memorial Sloan Kettering, and many others are accommodating a 15 percent average increase in volume, seeing wait times decreased by as much as 55 percent during peak hours, and reducing overtime hours by as much as 74 percent.

Put that in real-world terms: A one-hour wait becomes 27 minutes. Who wouldn't want a half-hour of waiting room time back? Time enough to even squeeze in one more game of Tetris. OI

Mohan Giridharadas is the founder and CEO of LeanTaaS, a Lean and predictive analytics company, Santa Clara, Calif.