

Machine-learning system determines the fewest, smallest doses that could still shrink brain tumors

August 10 2018, by Rob Matheson



MIT researchers aim to improve the quality of life for patients suffering from glioblastoma, the most aggressive form of brain cancer, with a machine-learning model that makes chemotherapy and radiotherapy dosing regimens less toxic but still as effective as human-designed regimens. Credit: Massachusetts Institute of Technology



MIT researchers are employing novel machine-learning techniques to improve the quality of life for patients by reducing toxic chemotherapy and radiotherapy dosing for glioblastoma, the most aggressive form of brain cancer.

Glioblastoma is a malignant tumor that appears in the brain or spinal cord, and prognosis for adults is no more than five years. Patients must endure a combination of radiation therapy and multiple drugs taken every month. Medical professionals generally administer maximum safe drug doses to shrink the tumor as much as possible. But these strong pharmaceuticals still cause debilitating side effects in patients.

In a paper being presented next week at the 2018 Machine Learning for Healthcare conference at Stanford University, MIT Media Lab researchers detail a model that could make dosing regimens less toxic but still effective. Powered by a "self-learning" machine-learning technique, the model looks at treatment regimens currently in use, and iteratively adjusts the doses. Eventually, it finds an optimal treatment plan, with the lowest possible potency and frequency of doses that should still reduce tumor sizes to a degree comparable to that of traditional regimens.

In simulated trials of 50 patients, the machine-learning model designed treatment cycles that reduced the potency to a quarter or half of nearly all the doses while maintaining the same tumor-shrinking potential. Many times, it skipped doses altogether, scheduling administrations only twice a year instead of monthly.

"We kept the goal, where we have to help patients by reducing tumor sizes but, at the same time, we want to make sure the quality of life—the dosing toxicity—doesn't lead to overwhelming sickness and harmful side effects," says Pratik Shah, a principal investigator at the Media Lab who supervised this research.



The paper's first author is Media Lab researcher Gregory Yauney.

Rewarding good choices

The researchers' model uses a technique called reinforced learning (RL), a method inspired by behavioral psychology, in which a model learns to favor certain behavior that leads to a desired outcome.

The technique comprises artificially intelligent "agents" that complete "actions" in an unpredictable, complex environment to reach a desired "outcome." Whenever it completes an action, the agent receives a "reward" or "penalty," depending on whether the action works toward the outcome. Then, the agent adjusts its actions accordingly to achieve that outcome.

Rewards and penalties are basically positive and negative numbers, say +1 or -1. Their values vary by the action taken, calculated by probability of succeeding or failing at the outcome, among other factors. The agent is essentially trying to numerically optimize all actions, based on reward and penalty values, to get to a maximum outcome score for a given task.

The approach was used to train the computer program DeepMind that in 2016 made headlines for beating one of the world's best human players in the game "Go." It's also used to train driverless cars in maneuvers, such as merging into traffic or parking, where the vehicle will practice over and over, adjusting its course, until it gets it right.

The researchers adapted an RL model for glioblastoma treatments that use a combination of the drugs temozolomide (TMZ) and procarbazine, lomustine, and vincristine (PVC), administered over weeks or months.

The model's agent combs through traditionally administered regimens. These regimens are based on protocols that have been used clinically for



decades and are based on animal testing and various clinical trials. Oncologists use these established protocols to predict how much doses to give patients based on weight.

As the model explores the regimen, at each planned dosing interval—say, once a month—it decides on one of several actions. It can, first, either initiate or withhold a dose. If it does administer, it then decides if the entire dose, or only a portion, is necessary. At each action, it pings another clinical model—often used to predict a tumor's change in size in response to treatments—to see if the action shrinks the mean tumor diameter. If it does, the model receives a reward.

However, the researchers also had to make sure the model doesn't just dish out a maximum number and potency of doses. Whenever the model chooses to administer all full doses, therefore, it gets penalized, so instead chooses fewer, smaller doses. "If all we want to do is reduce the mean tumor diameter, and let it take whatever actions it wants, it will administer drugs irresponsibly," Shah says. "Instead, we said, 'We need to reduce the harmful actions it takes to get to that outcome.'"

This represents an "unorthodox RL model, described in the paper for the first time," Shah says, that weighs potential negative consequences of actions (doses) against an outcome (tumor reduction). Traditional RL models work toward a single outcome, such as winning a game, and take any and all actions that maximize that outcome. On the other hand, the researchers' model, at each action, has flexibility to find a dose that doesn't necessarily solely maximize tumor reduction, but that strikes a perfect balance between maximum tumor reduction and low toxicity. This technique, he adds, has various medical and clinical trial applications, where actions for treating patients must be regulated to prevent harmful side effects.

Optimal regimens



The researchers trained the model on 50 simulated patients, randomly selected from a large database of <u>glioblastoma</u> patients who had previously undergone traditional treatments. For each patient, the model conducted about 20,000 trial-and-error test runs. Once training was complete, the model learned parameters for optimal regimens. When given new patients, the model used those parameters to formulate new regimens based on various constraints the researchers provided.

The researchers then tested the model on 50 new simulated patients and compared the results to those of a conventional regimen using both TMZ and PVC. When given no dosage penalty, the model designed nearly identical regimens to human experts. Given small and large dosing penalties, however, it substantially cut the doses' frequency and potency, while reducing tumor sizes.

The researchers also designed the model to treat each patient individually, as well as in a single cohort, and achieved similar results (medical data for each patient was available to the researchers). Traditionally, a same dosing regimen is applied to groups of patients, but differences in tumor size, medical histories, genetic profiles, and biomarkers can all change how a patient is treated. These variables are not considered during traditional clinical trial designs and other treatments, often leading to poor responses to therapy in large populations, Shah says.

"We said [to the <u>model</u>], 'Do you have to administer the same dose for all the patients? And it said, 'No. I can give a quarter dose to this person, half to this person, and maybe we skip a dose for this person.' That was the most exciting part of this work, where we are able to generate precision medicine-based treatments by conducting one-person trials using unorthodox machine-learning architectures," Shah says.

This story is republished courtesy of MIT News



(<u>web.mit.edu/newsoffice/</u>), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: Machine-learning system determines the fewest, smallest doses that could still shrink brain tumors (2018, August 10) retrieved 5 May 2024 from <u>https://medicalxpress.com/news/2018-08-machine-learning-fewest-smallest-doses-brain.html</u>

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.