


Cracking open the black box of automated machine learning

3 June 2019, by Rob Matheson



Researchers from MIT and elsewhere have developed an interactive tool that, for the first time, lets users see and control how increasingly popular automated machine-learning (AutoML) systems work. Credit: Chelsea Turner, MIT

Researchers from MIT and elsewhere have developed an interactive tool that, for the first time, lets users see and control how automated machine-learning systems work. The aim is to build confidence in these systems and find ways to improve them.

Designing a machine-learning model for a certain task—such as image classification, disease diagnoses, and stock market prediction—is an arduous, time-consuming process. Experts first choose from among many different algorithms to build the model around. Then, they manually tweak "hyperparameters"—which determine the model's overall structure—before the model starts training.

Recently developed automated machine-learning (AutoML) systems iteratively test and modify algorithms and those hyperparameters, and select the best-suited models. But the systems operate as "[black boxes](#)," meaning their selection

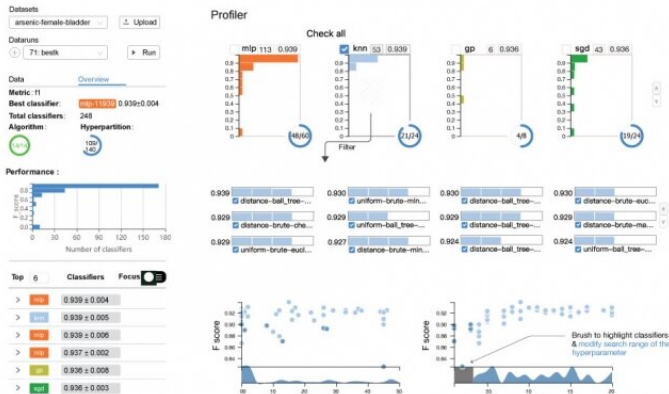
techniques are hidden from users. Therefore, users may not trust the results and can find it difficult to tailor the systems to their search needs.

In a paper presented at the ACM CHI Conference on Human Factors in Computing Systems, researchers from MIT, the Hong Kong University of Science and Technology (HKUST), and Zhejiang University describe a tool that puts the analyses and control of AutoML methods into users' hands. Called ATMSeer, the tool takes as input an AutoML system, a dataset, and some information about a user's task. Then, it visualizes the search process in a user-friendly interface, which presents in-depth information on the models' performance.

"We let users pick and see how the AutoML systems works," says co-author Kalyan Veeramachaneni, a principal research scientist in the MIT Laboratory for Information and Decision Systems (LIDS), who leads the Data to AI group. "You might simply choose the top-performing model, or you might have other considerations or use domain expertise to guide the system to search for some models over others."

In [case studies](#) with science graduate students, who were AutoML novices, the researchers found about 85 percent of participants who used ATMSeer were confident in the models selected by the system. Nearly all participants said using the tool made them comfortable enough to use AutoML systems in the future.

"We found people were more likely to use AutoML as a result of opening up that black box and seeing and controlling how the system operates," says Micah Smith, a graduate student in the Department of Electrical Engineering and Computer Science (EECS) and a researcher in LIDS.



The tool, ATMSeer, generates a user-friendly interface that shows in-depth information about a chosen models' performance, as well as the selection of algorithms and parameters that can all be adjusted. Credit: Massachusetts Institute of Technology

"Data visualization is an effective approach toward better collaboration between humans and machines. ATMSeer exemplifies this idea," says lead author Qianwen Wang of HKUST. "ATMSeer will mostly benefit machine-learning practitioners, regardless of their domain, [who] have a certain level of expertise. It can relieve the pain of manually selecting machine-learning algorithms and tuning hyperparameters."

Joining Smith, Veeramachaneni, and Wang on the paper are: Yao Ming, Qiaomu Shen, Dongyu Liu, and Huamin Qu, all of HKUST; and Zhihua Jin of Zhejiang University.

Tuning the model

At the core of the new tool is a custom AutoML system, called "Auto-Tuned Models" (ATM), developed by Veeramachaneni and other researchers in 2017. Unlike traditional AutoML systems, ATM fully catalogues all search results as it tries to fit models to data.

ATM takes as input any dataset and an encoded prediction task. The system randomly selects an algorithm class—such as neural networks, decision trees, random forest, and logistic regression—and

the model's hyperparameters, such as the size of a decision tree or the number of neural network layers.

Then, the system runs the model against the dataset, iteratively tunes the hyperparameters, and measures performance. It uses what it has learned about that model's performance to select another model, and so on. In the end, the system outputs several top-performing models for a task.

The trick is that each model can essentially be treated as one data point with a few variables: algorithm, hyperparameters, and performance. Building on that work, the researchers designed a system that plots the data points and variables on designated graphs and charts. From there, they developed a separate technique that also lets them reconfigure that data in real time. "The trick is that, with these tools, anything you can visualize, you can also modify," Smith says.

Similar visualization tools are tailored toward analyzing only one specific [machine-learning model](#), and allow limited customization of the search space. "Therefore, they offer limited support for the AutoML process, in which the configurations of many searched models need to be analyzed," Wang says. "In contrast, ATMSeer supports the analysis of machine-learning models generated with various algorithms."

User control and confidence

ATMSeer's interface consists of three parts. A control panel allows users to upload datasets and an AutoML system, and start or pause the search process. Below that is an overview panel that shows basic statistics—such as the number of algorithms and hyperparameters searched—and a "leaderboard" of top-performing models in descending order. "This might be the view you're most interested in if you're not an expert diving into the nitty gritty details," Veeramachaneni says.

ATMSeer includes an "AutoML Profiler," with panels containing in-depth information about the algorithms and hyperparameters, which can all be adjusted. One panel represents all algorithm classes as histograms—a bar chart that shows the

distribution of the algorithm's performance scores, on a scale of 0 to 10, depending on their hyperparameters. A separate panel displays scatter plots that visualize the tradeoffs in performance for different hyperparameters and algorithm classes.

Case studies with machine-learning experts, who had no AutoML experience, revealed that user control does help improve the performance and efficiency of AutoML selection. User studies with 13 graduate students in diverse scientific fields—such as biology and finance—were also revealing. Results indicate three major factors—number of algorithms searched, [system](#) runtime, and finding the top-performing [model](#)—determined how users customized their AutoML searches. That information can be used to tailor the systems to users, the researchers say.

"We are just starting to see the beginning of the different ways people use these systems and make selections," Veeramachaneni says. "That's because now that this information is all in one place, and people can see what's going on behind the scenes and have the power to control it."

More information: Qianwen Wang et al.

ATMSeer, *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19* (2019). [DOI: 10.1145/3290605.3300911](#)

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

APA citation: Cracking open the black box of automated machine learning (2019, June 3) retrieved 1 October 2022 from <https://techxplore.com/news/2019-06-black-automated-machine.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.