


Machine learning

This cutting-edge technology could revolutionize GA maintenance

BY MIKE BUSCH



THE EXHAUST VALVE is the most likely component of a piston aircraft engine to fail catastrophically. When one fails, combustion ceases in the cylinder, and the engine loses power and starts running rough. This usually results in a precautionary landing—on-airport if you’re lucky, off-airport if you’re not. It’s particularly serious with a four-cylinder engine, because a four runs a lot worse on three than a six does on five. Occasionally, the liberated valve fragment gets wedged into the piston crown, which can shatter the piston and cause a total power loss, sometimes with fatal results.

That’s why it’s so important to detect incipient exhaust valve failures early before they can cause serious problems. The traditional way is the annual compression test, but that isn’t a very reliable method because loss of compression typically doesn’t occur until the valve is very sick and close to failing.

A borescope inspection is vastly better and can detect failing valves a lot earlier. The AOPA Air Safety Institute publishes a great poster to help mechanics understand what to look for—Google “Anatomy of a Valve Failure.” Unfortunately, most shops don’t do regular borescope inspections, and the ones that do typically do them only at the annual. That may not be often enough, especially for engines that fly a lot.

There should be a better way. If your airplane has a recording digital engine monitor—and nowadays over half the piston fleet does—the answer might be hidden in your engine monitor data.

Predictive analytics

Predictive analytics is the use of data, statistical algorithms, and machine learning techniques to identify the likelihood of future outcomes based on historical data. Instead of just using data to analyze what happened, the goal of predictive analytics is to assess what is likely to happen in the future.

Over the past decade, remarkable advances in this technology have been made in the airline industry. Boeing has been in the forefront with its groundbreaking Airplane Health Management program for the 777 and 787, but Airbus, GE, Pratt & Whitney, and Rolls-Royce have been investing heavily in this area. The resulting capability to use aircraft sensor data to predict incipient component failures has been remarkable.

Piston GA ought to be able to benefit from this technology, too, I thought, so in 2014 my company launched a program whose goal was to predict incipient exhaust valve failures based on digital engine monitor data. We dubbed it FEVA—Failing Exhaust Valve Analytics.

FINALLY!
A CARD
FOR PILOTS.



Cash Back¹
Every Purchase

A CASH BACK¹ REDEMPTION IS APPLIED AS A STATEMENT CREDIT.
QUARTERLY BONUS POINT CAP OF 2,500 POINTS

Switch today to the AOPA World Mastercard[®], and stop paying too much on your aviation purchases.

Visit
AOPA.org/creditcard
to learn more.



*Certain points and purchases restrictions apply, see full Rewards Terms and Conditions for full details at AOPA.org/creditcard.
1. Rewards points can be redeemed for Cash Back or other items provided through AOPA Pilot Rewards. A Cash Back redemption is applied as a statement credit. The statement credit will reduce your balance but you are still required to make at least your minimum payment. A minimum of 2,500 points is needed to redeem for Cash Back. Values for non-cash back redemption items such as merchandise, gift cards, and travel may vary.

Humble beginnings

Over many years of looking at charted digital engine monitor data from piston aircraft engines, I occasionally observed an unusual pattern—a slow, rhythmic exhaust gas temperature oscillation with a frequency of roughly one cycle per minute—that seemed to correlate strongly with burned exhaust valves. I trained my staff of analysts to keep an eye out for this pattern, and to alert clients whenever they saw it so that the corresponding cylinder could be borescoped. Usually, cylinders exhibiting this peculiar EGT oscillation turned out to have failing exhaust valves.

Data from more than 3,000,000 GA flights has now been uploaded to our charting platform, but only a tiny fraction of those have been examined by our human analysts. I decided we should create a computer algorithm—what we now call FEVA 1.0—to scan every flight uploaded to the platform for this EGT oscillation. Whenever the algorithm spotted the pattern (as it did in EGT #5 in the illustration), the software would alert a human analyst to look at the flight. If the analyst agreed that the EGT pattern looked suspicious, we'd alert the aircraft owner and suggest a borescope inspection of the cylinder.

FEVA 1.0 was an “expert system”—an algorithm that attempts to emulate the decision-making ability of a human expert. We back-tested it on hundreds of actual flights, tweaking the algorithm to get it to alert on the same flights as our analysts did, and fine-tuning it to achieve optimum balance between maximum detection sensitivity and minimum false-positive rate.

Not good enough

FEVA 1.0 was a step in the right direction, and on more than 60 occasions predicted exhaust valve failures before they happened. But it wasn't good enough. Clients would sometimes send us borescope images of valves that were obviously extremely sick and close to failure, and ask “why didn't FEVA predict this?” That's embarrassing.

FEVA 1.0 was just too narrow-minded. Cylinders exhibiting slow rhythmic EGT oscillation usually have a failing exhaust valve, but the converse is not true.

4 STEPS TO PURCHASING YOUR FIRST AIRCRAFT

1



Understand your financial picture.

2



Apply and get approved.

3




Negotiate the purchase.

4



Close on your first Aircraft.



Call for a quote today!

1.844.674.2761

AOPAFinance.org/FirstTimeBuyer



AOPA FINANCE

Cylinders with failing exhaust valves don't always exhibit these telltale EGT oscillations. If the valve's rotator cap stops working properly or if the valve starts sticking in its guide because of deposit buildup, the valve might start heating unevenly and burning badly, but the EGT will not oscillate. FEVA 1.0 couldn't detect such failures, and our human analysts might not spot them, either.

It became apparent that any algorithm capable of reliability predicting incipient exhaust valve failures might need to look for more than just EGT oscillations and might need to look at more than just EGT. It might need to look for complex patterns and interrelationships among data obtained from many different sensors. What sensor data should it be looking at? What patterns and interrelationships should it be looking for? We had no clue. Nor, for that matter, did Continental, Lycoming, Rotax, or anyone else in the industry.


"I bet we could get the computer to figure this out!" exclaimed my colleague Chris Wrather to me over brunch one day. Chris has a unique background—entrepreneur, airplane owner, instrument pilot, A&P mechanic, computer programmer, and Ph.D.-level

of input variables than humans are. Also, machine learning models keep getting smarter over time as they get fed more training data.

Chris became convinced that a FEVA 2.0 based on machine learning was a better way to go. His plan was to program the computer with a generalized machine-learning algorithm known as the "random forests" model whose inputs would include pretty much everything we could get from the aircraft's digital engine monitor—more than 30 different input variables.

The model would then be "trained" by feeding it data from thousands of actual GA flights, some involving engines known to have failing exhaust valves, others from engines with valves known to be healthy. For each flight in this "training set" we would tell the model whether the flight was from a healthy-valve engine or a sick-valve engine and which cylinder(s) had failing valve(s). It would then be up to the machine learning model to figure out how to tell the difference.

After much training and back-testing, we deployed FEVA 2.0 about a year ago. It proved superior to FEVA 1.0 in every respect. It caught more failing exhaust valves. It generated fewer false alarms.



The slow, rhythmic oscillation of EGT #5 suggests that the exhaust valve is failing.

mathematician. Chris had just completed an online course in machine learning offered by Stanford University.

A better mousetrap

Suppose you wanted to create a computer algorithm that could analyze a digital photo and determine whether it was of a man or a woman. One way of doing that would be to assemble a group of experts and ask them to come up with a set of rules—such as men usually have broader shoulders and more facial hair, women have more curves and longer hair—and programming those rules into the computer. Another way would be to feed the computer thousands of photos, each one identified as male or female, and let the computer figure out for itself how to tell which is which. The former is called an "expert system" while the latter is called "machine learning."

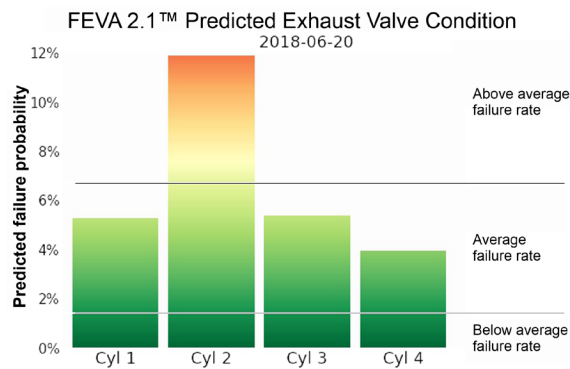
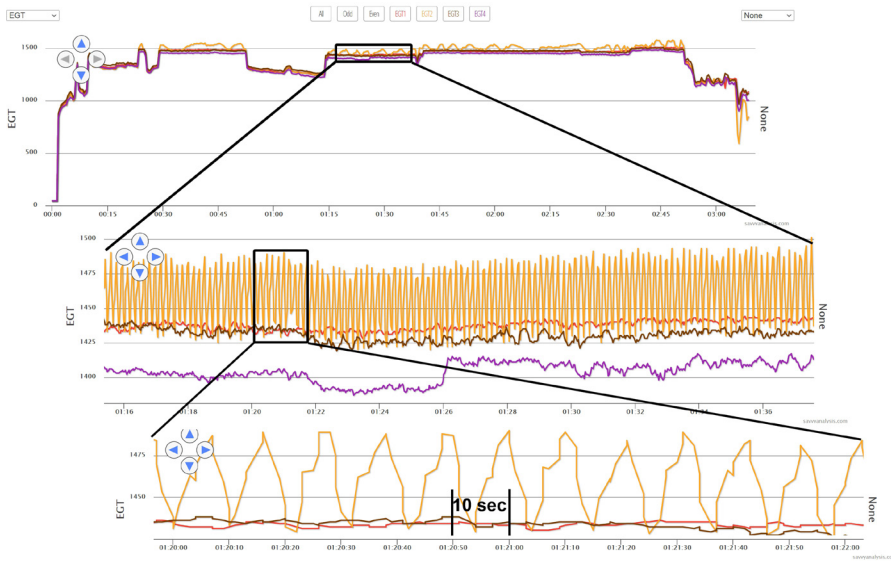
Machine learning has some distinct advantages over the expert system approach. Computers are generally a lot better at figuring out complex patterns and interrelationships among a large number

In contrast to FEVA 1.0's simple go/no-go output, the FEVA 2.0 model yielded a numeric "risk score" for each cylinder that could range anywhere from very low risk to very high risk. Any time the model produced an above-average risk score for one or more cylinders, we would advise the aircraft owner to have those cylinders borescoped ASAP.

In May, we released FEVA 2.1, a significantly improved version of the machine learning. This improved version differs from its predecessor in two important ways:

- It extends the list of input variables to more than 35; and
- It draws its input data from an aircraft's entire recent flight history, rather than from only the most recent flight as was previously the case.

The new model is both 20 percent better at predicting failing valves ("sensitivity") and 20 percent less likely to make a false prediction ("positive predictive value") than its predecessor. In addition, successive reports will be more consistent over time.



The proof is in the pudding

To illustrate how far we’ve come in the past seven years, a case study may be instructive. In 2015, Wally purchased a beautiful 2008 DA40 Diamond Star. Wally has been religious about uploading the data captured by his Garmin G1000 to the SavvyAnalysis platform. In July 2018, Wally put his airplane in the shop for its annual inspection. During the compression test on the airplane’s Lycoming O-360, the compression on cylinder number 2 was 0/80. A borescope inspection revealed a seriously burned exhaust valve in that cylinder that probably would have failed soon had it not been caught at the annual.

Leading up to the annual, each of Wally’s flights had been scanned by FEVA 1.0, and none raised a red flag that the number 2 exhaust valve might be failing. Why not?

In 2018, FEVA was still an expert system programmed to alert whenever it found a slow, rhythmic EGT oscillation that we knew was reliably indicative of a failing exhaust valve. If you look at the EGT data from Wally’s last flight before the annual inspection (top), you can see that EGT number 2 looks wonky.

But it’s not the sort of slow, rhythmic oscillation that FEVA 1.0 was programmed to look for, but rather a much faster rhythmic oscillation. In fact, EGT number 2 was oscillating at a

frequency of about one cycle per 10 or 11 seconds, six times faster than the one cycle per minute oscillation that FEVA 1.0 was designed to look for. In fact, FEVA 1.0 was specifically programmed to ignore EGT oscillations faster than one cycle per 30 seconds based on the “expert assumption” that oscillations faster than that almost had to be noise. No wonder it didn’t alert on Wally’s failing exhaust valve.

So much for the expert. (The expert was me. Sigh.)

Would the machine learning model have done better? We wondered the same thing, so we ran Wally’s flights prior to the 2018 annual inspection through FEVA 2.1’s machine learning model. The results are at left.

Wow! The machine learning model returned an off-the-charts risk score for cylinder 2, one of the highest risk scores we’ve ever seen the model produce. How exactly did it conclude that the number 2 cylinder was at extraordinarily high risk for exhaust valve failure? That’s hard to say.

The model uses more than 35 input variables to make its predictions. It looks for hidden patterns in the data, some of which might not be apparent to the human eye. While the machine learning approach we use is enormously powerful, one drawback is that it is difficult to tease out the effects of individual variables to arrive at an explanation of why the model made the prediction that it did.

What we can say is that the model did not have access to information from the compression test or the borescope inspection. It looked solely at engine monitor data, and it looked over a series of flights in order to come up with its prediction that the number 2 exhaust valve was at high risk of failing.

This is exciting technology, and I think we’ve just scratched the surface of what it can do for piston GA maintenance. My team is working hard to make predictive analytics smarter and more accurate, and to extend it to other safety-critical areas such as sticking valves, failing ignition, and detonation/preignition. If you’re flying a piston aircraft, predictive analytics might save your bacon someday. mike.busch@savvyaviation.com ► savvyaviation.com

Ask the A&Ps /

Every month, experts Mike Busch, Paul New, and Colleen Sterling answer your toughest aviation maintenance questions on our *Ask the A&Ps* podcast. Submit your questions to podcasts@aopa.org.

► aopa.org/news-and-media/podcasts/podcasts/ask-the-a-and-ps