

Machine Learning: Using Algorithms to Sort Fruit

At Amazon, we've been engaged in the practical application of machine learning for many years now. Some of this work is highly visible, but much of what we do with machine learning happens beneath the surface – quietly but meaningfully improving our core operations.

By About Amazon Staff

Over past decades computers have broadly automated tasks that programmers could describe with clear rules and algorithms. Modern machine learning techniques now allow us to do the same for tasks where describing the precise rules is much harder.

At Amazon, we've been engaged in the practical application of machine learning for many years now. Some of this work is highly visible: our autonomous Prime Air delivery drones; the Amazon Go convenience store that uses machine vision to eliminate checkout lines; and Alexa, our cloud-based AI assistant. But much of what we do with machine learning happens beneath the surface. Machine learning drives our algorithms for demand forecasting, product search ranking, product and deals recommendations, merchandising placements, fraud detection, translations, and much more. Though less visible, much of the impact of machine learning will be of this type – quietly but meaningfully improving core operations.

For example, we have developed the world's first automated ripeness detection system for fruit and vegetables. The unique technology uses machine learning and algorithms to determine the quality of the produce, thus promising more reliability for online grocery shopping for customers, and a significant reduction in food waste. The problem of good fruit being mistakenly thrown away, or expired produce being delivered to the customer, will be a thing of the past.

In most AmazonFresh depots, ripeness is still checked by hand but assessing the ripeness levels accurately is difficult even for professional inspectors. As our tests showed, assessors arrived at different results in 20% of the tests despite checking the quality of the exact same piece of fruit. This is why our team of machine learning experts and the AmazonFresh team joined forces to improve automation in the quality control of fresh produce.

The automated ripeness detection system consists of a conveyor belt that transports the food in containers to a particular sensor. The sensor looks like a normal camera, but it can capture information which is invisible to the human eye. We teach the machine what good and bad products look like by inputting new product variants on a daily basis. The products are photographed and made available to the machine in the shape of data. That way, the computer gradually understands the quality standards.

The fresh produce is divided into four categories: "OK", "Damaged", "Badly damaged" and "Expired". Containers of the different ripeness categories are randomly offered to the machine. Likewise, employees only find out via a screen with which goods they are supposed to fill the entire container. This is to avoid every third container being filled with "expired" produce, or every fourth container with "OK" fruit. It means that rather than allowing the machine to learn a sequence instead of conducting proper product tests, we actually prevent the machine from learning a pattern.

Love this article?

Get many more like it, delivered right to your inbox.

Protected by reCAPTCHA. The Google [Privacy Policy](#) and [Terms of Service](#) apply.

[Amazon Privacy Policy](#) | Opt out anytime

The freshness conveyor belt is still in the early stages of development and needs further improvements over the coming years. Until then, many challenges need to be solved. During a test run in the US, the system promptly ground to a halt. The reason: Although developed by international teams, the system does not (yet) like American apples. They are much too large and caused blockages and traffic jams on the belt during the test run.

Even more exciting insights will be possible in the future. Michael Donoser, team lead of the Computer Vision department at Amazon Machine Learning reveals: "We are constantly developing the algorithms further. The next goal is for the machine to not only determine the ripeness of a piece of fruit, but also whether it tastes sweet or sour."