

# Machines That Learn

Created: November 2017

Last Updated: March 2018

Estimated time:	55 minutes <ul style="list-style-type: none"><li>• [30 minutes] Activity #1</li><li>• [25 minutes] Assignment</li></ul>
Group or individual activity:	Individual
Ages:	13-18 years old
Grades:	Grades 8-12
Online/offline elements:	This learning experience contains links to online resources, with an offline writing assignment.
Areas:	Main area: Artificial Intelligence Additional areas: Computational Thinking, Data Literacy, Information Literacy
License:	<a href="https://creativecommons.org/licenses/by-sa/4.0/">Creative Commons Attribution-ShareAlike 4.0 International</a> license. For more information, please visit: <a href="http://dlrp.berkman.harvard.edu/about">http://dlrp.berkman.harvard.edu/about</a>

## Learning Goal

Gain a basic understanding of machine learning, its forms, and its limitations.

## Materials

- Paper

- Pen or pencil
- [Optional] Computer or mobile device

## Resources

- Video: [What Is Machine Learning?](#) - by Android Authority
- Visualization: [A Visual Introduction to Machine Learning](#) - by R2D3

## Activity #1: How Do Machines Learn?

### Introduction

Patterns help us make many of our daily decisions. For instance, if you are going to the movies, and a new super hero film has just come out, you may need to decide whether or not to see it. You might decide based on something like whether or not you enjoyed other super hero movies or if you like comic books. If this type of information helped you make a decision, you were basing your decision on previous data.

You could also use a machine learning (ML) algorithm to predict whether or not you would like a certain movie. The difference is that a learning algorithm could go through [many more examples](#) (or data) about which movies you like than you could go through in your head.

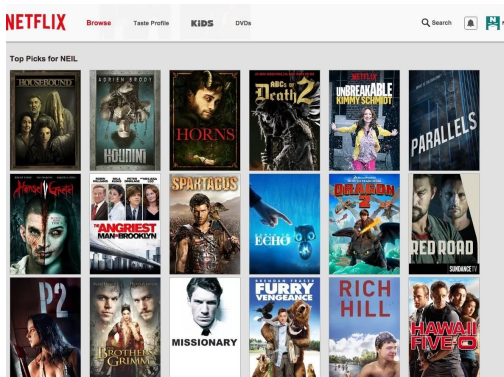


Figure 1) A [Netflix user's home page](#), generated by machine learning

This is actually how [Netflix's algorithm](#) works. It takes information about every movie you have watched and the ratings you have given, and then it finds new movies you might like based on the fact that the movies share similar characteristics.

Machine learning is a subdiscipline of artificial intelligence, which relies on large datasets and using programs that perform tasks without having to be explicitly told what to do, unlike [older AI techniques like rule-based systems](#). This allows us to use these algorithms to predict future outcomes based on current or past information.

## Types of Machine Learning

Machine learning helps us to see patterns, but ML algorithms can find patterns in several ways. Here are a few of them:

**Supervised learning** refers to algorithms that learn from data that we give them. For example, if you want a computer to be able to detect images of lizards, you could give it thousands of images of lizards.



Figure 2) A photo of an [iguana](#)

**Unsupervised learning**, as you might have guessed, refers to algorithms that don't require any labeled data. With the same lizard example, an unsupervised learning algorithm could take thousands of images from the results of a Google image search and would group them according to whatever patterns were most meaningful. This could potentially help us learn new things about the images because the algorithm can show us patterns we didn't know about.

Finally, **reinforcement learning** is a type of algorithm that is able to be ‘reinforced’ by how well the learning algorithm achieves an abstract goal in its environment. For instance, the algorithm in a self-driving car would have the abstract goal of driving its passengers safely and effectively to their destination. So, it would be given a “reward” feedback when it correctly drives, brakes, turns, etc. This helps the computer ‘learn’ which actions result in a reward (like stopping at a red light) and which do not (hitting a mailbox), so it can shape its future behavior to increase its rewards.

There are many different forms of machine learning algorithms, which often require various types or amounts of data, algorithms, and computing resources. Some of these algorithms might be better suited for specific situations, such as different types of datasets (e.g., datasets for large businesses). If you want to explore these types further, you can watch this [video](#).

### **Criticisms and Limitations**

Machine learning allows us to do more with data than ever before (as you’ve probably heard about), but it does come with concerns and limitations.

One concern is that these models require a very large amount of data, which can be expensive or hard to acquire. This creates a market for big data, which incentivizes companies like Google and Facebook to use their services to collect and sell user data.

One way this data might be used is by advertisers running it through machine learning algorithms in order to help them more successfully target ads at consumers than was possible with other techniques. This can result in many more sales of a product, so it is a valuable tool. In fact, these days, data itself is becoming more of a business asset than the software powering these big companies.

Another problematic issue in machine learning is known as “overfitting.” Overfitting is when a machine learning algorithm performs well on the data that was used to train it, but poorly when it encounters new data. This could mean that the pattern identified by the machine learning algorithm works well with the training data, but is inaccurate or nonexistent in reality.

For instance, some colleges have a “legacy program” that gives preference to new applicants who have had a family member who graduated from their institution. Now, imagine that the university ended this program but was using an ML algorithm to analyze the chance of a new student getting accepted. The algorithm would predict that

having a family member who had graduated from the school would increase the chances of being accepted, despite the fact that this was no longer true. The machine learning program 'overfit' the data because it assumed that a pattern that was true in the test data (past acceptances) would be true in future data (current acceptances). This [article](#) explains such limitations in more depth.

## Assignment

Come up with a task that could be completed by a digital device (e.g., desktop computer, mobile device) using a machine learning algorithm. In a paragraph or two (written or typed), describe the types of datasets you think you would need to train the algorithm. Identify the type of learning algorithm that would be best to use (supervised, unsupervised, or reinforcement).

### Example:

I take a lot of pictures of my dogs with my phone. I want it to automatically add these pictures to a specific folder, so I don't have to scroll through all of my photos to find pictures of my dogs. In order to do this, I would need to train a machine learning algorithm with lots of pictures of my dogs that are labeled as 'my dogs' as well as pictures of things that are not my dogs that are labeled as 'not my dogs.' This type of algorithm is known as supervised learning.



Figure 3) A photo of [two white dogs](#)