Cloud Machine Learning Services (https://cloud.google.com/products/machine-learning/)

What is machine learning?

Machine learning is functionality that helps software perform a task without explicit programming or rules. Traditionally considered a subcategory of artificial intelligence, machine learning involves statistical techniques, such as *deep learning* (aka neural networks), that are inspired by theories about how the human brain processes information.

How does machine learning work?

Machine learning is powered by algorithmic models that are trained to recognize patterns in collected data (such as logs, speech, text, or images). Because access to lots of training data and computing power are preconditions for success, the cloud (where data storage and high-performance computation are plentiful and can be particularly cost-efficient) is an ideal platform for machine learning.

What are some examples of machine learning?

Examples of machine learning abound in everyday experiences. A very simple example would be the auto-completion of names, keywords, or addresses in a search field, but the same concept can be applied in more complex use cases across multiple industries. For instance, machine learning is used to:

- Categorize images (such as MRI studies, photos, or satellite imagery)
- Look for keywords in massive numbers of text documents or emails
- Flag potentially fraudulent transactions
- Personalize product recommendations based on customer behavior

- Enable software to accurately respond to voice commands
- Predict weather patterns or other climate conditions
- Translate languages in text or audio

In summary, wherever there is software that performs a labor-intensive task on a scale beyond human capability, machine learning may well be involved.

How is machine learning used inside Google?

Machine learning has been a cornerstone of Google's internal systems for years, primarily because of our need to automate data-driven systems on a massive scale. This experience has provided unique insight into the right frameworks, techniques, infrastructure, and data that can help customers complete a successful journey toward getting value out of machine learning.

For example, the open source framework called TensorFlow, originally developed for use inside Google, is now a standard in the data science community, giving advanced users the ability to build and train powerful models. And, in addition to heavily contributing to the academic and open source communities, our hundreds of machine-learning researchers help bring that functionality into Google products (such as G Suite, Search, and Photos) in addition to Google's internal operations (for data center automation, for example).

How can I start using machine learning in my own organization?

Once belonging to the exotic domains of statistics and data science, machine-learning capability is now widely accessible in the form of open source libraries (TensorFlow), as well as managed services and cloud APIs. For data scientists who want to build "future-proof" models that can move between on-premises and the cloud, or mainstream developers who lack adequate training data and want to bring a pre-built/pre-trained model into their app via a cloud API, using such tools as part of the daily routine is a realistic goal.

Learn more:

- Google Machine Learning (https://cloud.google.com/products/machine-learning/)

- Predictive Analytics (https://cloud.google.com/ml-engine/)
- Data and Machine Learning Training (https://cloud.google.com/training/data-ml)

Last updated September 27, 2019.