Xiaoxiang Zhu

Watching **Cities Grow**

Highly precise computational models of the world's mega-cities could help to improve urban planning. Working towards this goal, Xiaoxiang Zhu is combining petabytes of earth observation satellite data and social media data for the first time.

Bjoern Menze

Getting a **Diagnosis from Pixels**

Bjoern Menze is training computer programs to "learn" how to detect brain tumors and propose individualized treatments.

Machine **Learning**

A lot of scientists are using machine learning methods in order to unveil information hidden in huge amounts of data. Remote sensing (p. 28) and medical image recognition (p. 38) are two examples highlighted in this issue. How does machine learning work and what impact does it have in research? We explore the key points.

What is machine learning?

Machine learning involves a computer autonomously developing knowledge from experience – similar to human learning – and independently coming up with solutions to new and previously unknown problems. To enable experience-based learning, a computer program analyzes examples and uses algorithms in an attempt to identify specific patterns and rules in the data. These are then applied to new cases.

How does it work?

Machine learning works along the same lines as human learning. Similar to the way a child learns to recognize certain objects in pictures, for instance, a computer can also "learn" to identify objects or distinguish between people. To enable this, a learning software program must first be populated with data and trained. Programmers will tell the system that a specific object is "a dog" and another "not a dog", for example. Next, the learning software receives continual feedback from the programmers, using the algorithm to adapt and optimize the model. Improving with each new data set, the model is ultimately able to make a clear distinction between dogs and non-dogs.

What is deep learning?

Deep learning is the most successful machine learning technique to date. Its methods are based on artificial neural networks. Following the example of the human brain, these simulate a network of connected neurons. Learning from experience is achieved by changing the strength of the connections between simulated neurons. In this way, machines can acquire abilities such as seeing, hearing, speaking, reading and writing. Deep learning eliminates many stages of conventional neural networks, since the computer covers all the intermediate steps.

Why machine learning and why now?

Machine learning stands to benefit both science and business. In science, it can be used to identify patterns and derive rules from data, for instance. Fields such as text analysis, machine translation and speech and image recognition are beginning to undergo radical changes thanks to machine learning.

Machine learning was originally based on research in the field of pattern recognition, carried out as far back as the 1980s. Technical constraints meant the area then lay dormant for quite a long time. However, just a few years ago, machine learning experienced a breakthrough in the form of parallel data processing capabilities, enabled by graphics processing units (GPUs) originally developed for the video gaming industry. GPUs are equipped with thousands of processing cores, making them significantly faster than standard processors.

Further developments such as multi-core architectures, improved algorithms and ultrafast in-memory databases have also had a positive impact. Another substantial factor is the growing availability of large amounts of structured and unstructured data from a wide range of sources – from sensors to digitalized documents and images – which can be used to "train" learning algorithms.

Klaus Manhart